

Nonlinear Dynamics in Distributed Systems

I. Adjali, J.L. Fernández-Villacañas and M. Gell

*Systems Research Division, BT Laboratories
Martlesham Heath, Ipswich IP5 7RE
United Kingdom
iadjali@bt-sys.bt.co.uk*

Abstract

We build on a previous statistical model for distributed systems and formulate it in a way that the deterministic and stochastic processes within the system are clearly separable. We show how internal fluctuations can be analysed in a systematic way using Van Kampen's expansion method for Markov processes. We present some results for both stationary and time-dependent states. Our approach allows the effect of fluctuations to be explored, particularly in finite systems where such processes assume increasing importance.

1 Introduction

With the increasing complexity of telecommunication and computational systems, an urgent requirement is developing for theoretical frameworks for addressing basic principles of distributed systems [1]. At present there is insufficient understanding of principles required to predict performance, to explain behaviour and to establish design methodologies [2]. The substantial vacuum in theoretical bases for distributed communication and computational systems stems largely from the historical preoccupation of computer and telecommunication science with uniprocessor systems [3].

The emergence of large decentralised systems is giving rise to the need for a general theoretic guide to the behaviour of large collections of locally controlled, asynchronous and concurrent processes interacting with an unpredictable environment. In particular this requires understanding the relation between the overall behaviour of the distributed system and that of its constituents, whose decisions are based upon local, imperfect, delayed and conflicting information. In many other systems, particularly in nature and societies, distributed systems with very complex behaviour and modes of operation have evolved. There is a growing awareness that many of the theoretical tools which have been developed with considerable success to describe distributed systems in physics [4], particularly in condensed matter physics, may be exploited in other fields, such as biology and economics [5] [6] [7] [8].

2 The Model

A central feature of open systems is the non-linear nature of their dynamics, which gives rise to a rich repertoire of behavioural regimes ranging from stable equilibrium to oscillations and chaotic states. One has to construct a model which integrates both the deterministic evolution equation, responsible for the macroscopic behaviour of the system, and the stochastic part which deals with fluctuations within the system.

A model was formulated for describing a self-organising open computational system with resources, free agents and pay-off mediated interactions [9] that builds on as well as overcomes the limitations inherent in previous work [10] [11].

Our approach allows the effects of fluctuations to be investigated systematically in the form of a large-system size expansion due to Van Kampen [12] [13]. Figure 1 shows an outline of the model; after writing down a probabilistic evolution equation for a general agent-resource system, we interpret the master equation obtained as describing a Markovian jump process and go on to apply Van Kampen's system size expansion. The deterministic equation for the behaviour of the system arises as the lowest-order term in the expansion and coincides with the mean-field equation of Kephart *et al.* [11]. The main contribution of the fluctuations comes in the form of a linear Fokker-Planck equation (FPE). Up to this order the noise in the system is linear and the solution of the master equation is given by a Gaussian distribution. Non-linear effects of fluctuations are calculated as small perturbations to the linear noise approximation. A detailed derivation of the equations in Fig. 1 can be found in Ref. [9]. In the following section some results for a two resource system in both stationary and time-dependent states are presented and discussed. Section 4 summarises the main points and indicates directions for further work.

3 Results and discussion

Our main objective is to investigate the approximation scheme for fluctuations based on the large system size expansion of Van Kampen. The one-step Markovian formulation of the problem allows us, in particular, to calculate the exact probability distribution for time-independent solutions (see eq. (12) in [9]). We shall begin by restricting our numerical calculations to stationary solutions of the system in order to make a direct comparison with exact results and test the validity of the approximation over a range of parameter values.

The function ρ represents the probability that an agent in the system will find resource 1 to be more attractive than resource 2. In general, the exact form of ρ is not known and will depend on several features of the problem at hand, such as incomplete, uncertain or delayed information about the available resources. As a first example, we make ρ a function of the payoffs G_1 and G_2 for using resources 1 and 2 respectively (as in [11]),

$$G_1 = 7 - f_1 \quad \text{and} \quad G_2 = 7 - 3f_2 \quad (1)$$

These pay-off functions model a simple competitive behaviour (opposing gradients) between agents so that the payoff for using each resource decreases with the number of agents already using the same resource. The system reaches a stability point when the two pay-offs are equal so agents will prefer staying with the resource they are using. For G_1 and G_2 given in (1) this optimal behaviour of the system




Figure 1: Main steps in the analysis of the proposed model for an agent-resource system

occurs for $f = 0.75$, that is, 75% of all agents using resource 1. The decision region can be made less sharply defined by introducing an uncertainty element in the payoff evaluation. This can be achieved by introducing Gaussian noise with standard deviation σ around the true value of the pay-off. If we assume that the agents' perception of each resource is different, then there will be one uncertainty parameter for each resource, σ_1 and σ_2 . The resulting transition probability ρ is given by,

$$\rho = \frac{1}{2} \left[1 + \operatorname{erf} \left(\frac{G_1 - G_2}{\sqrt{2} \sqrt{\sigma_1^2 + \sigma_2^2}} \right) \right] \quad (2)$$

and shown in Figure 2 for $\sigma_1 = \sigma_2 = 0.125$. The two limiting cases of $\sigma_1 = \sigma_2 = 0$ and $\sigma_{1,2} = \infty$ correspond respectively to perfect knowledge ($f = 0.75$) and complete lack of information on pay-offs, leading to the uniform distribution of agents ($f = 0.5$). By approximating f with its deterministic contribution ϕ , we obtain a graphical solution of the deterministic equation ($\rho(\phi) = \phi$) represented in Figure 2 by the crossing point between the curves $\rho(\phi)$ and ϕ . This point gives the equilibrium solution which now, due to a non-zero value of the uncertainty, is slightly offset from the optimal value $f = 0.75$. The macroscopic value of f for $\sigma_1 = \sigma_2 = 0.125$ is $\phi = 0.724$.

In order to see how the Van Kampen approximation depends on the uncertainty parameters, we have plotted the time-independent probability distributions for different orders in the approximation as well as the exact distribution, for three values




Figure 2: The transition probability (from resource 2 to resource 1) $\rho(f)$ corresponding to the pay-offs given in equ. (1), for two values of the uncertainty parameter $\sigma = \sigma_1 = \sigma_2$. The intersection with the line $\rho = f$ gives the solution to the time-independent macroscopic equation.

of $\sigma = \sigma_1 = \sigma_2$ (vertically) and three values of the number of agents N (horizontally) in Figure 3. We can draw the following conclusions: The approximation works reasonably well for all values of σ considered; the first order non-linear corrections are sufficient for correctly estimating fluctuation effects in the system, especially if the uncertainty parameter is not too small. Furthermore it seems that the approximation is best suited for systems with a moderate value of the uncertainty parameter ($\sigma \approx 0.5$), where non-linear effects of fluctuations, although significant, converge rapidly in the expansion. This may be the range of σ to look for in realistic systems, where agents are neither expected to have perfect knowledge nor be completely ignorant about the pay-offs of their transactions.

So far in our analysis we have only looked at systems with a single macroscopic stable behaviour, a consequence of the unique (stable) fixed point occurring at the intersection between the linear pay-off functions G_1 and G_2 . This simple competitive behaviour can be changed by making the pay-off functions non-linear, ie. introducing cooperation as well as competition between agents in the system. Whereas competition meant that agents would favour a resource if it had less agents using it, cooperation is expressed by an increased pay-off when a resource is used by more agents. The interplay of these two tendencies through non-linear pay-offs leads to a richer range of possible behaviours in the system. We treat here the example of a bistable system (arising from cubic pay-offs) (see Figure 4). Depending on which side of the mid-point the initial distribution is, the system will eventually settle in one of the macroscopic states characterised by the two peaks in the time-independent probability distribution.

The system's dynamics depends notably on the different values of the uncertainty parameters σ_1 and σ_2 . In Figure 4 we have shown the dependence on σ_1 for four different values of σ_2 , using the non-linear (cubic) pay-offs. By increasing σ_1 , σ_2 or both, the two peaks are seen to gradually get closer to each other and merge into a single (symmetric) peak. These critical values of σ_1 and σ_2 can be found by inspection of the time-independent macroscopic equation; they play the role

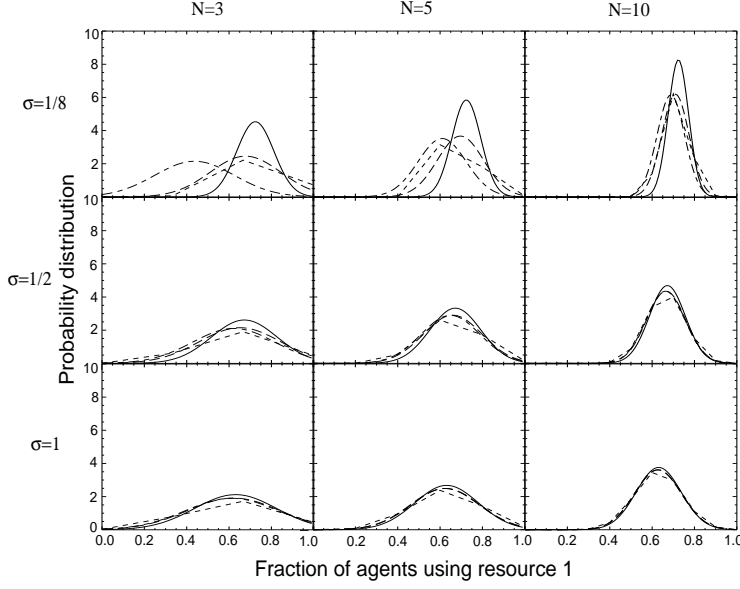


Figure 3: The time-independent probability distribution for three different values of N and three different values of $\sigma = \sigma_1 = \sigma_2$, in three orders of the large system-size expansion; full line: mean-field (including linear noise) result, long dash: first order non-linear corrections included, long dash-short dash: second order non-linear corrections included, short dash: exact solution.

of control parameters which can change qualitatively the dynamical phase space of the system, in this case from a system with two attractors to a system with a single one. This is reminiscent to phase transitions in physical systems such as the spontaneous magnetisation of a ferromagnetic system which happens by lowering the temperature below a critical value (Curie temperature). Above this value the overall magnetisation is zero and symmetric while below it there are two possible states of opposite magnetisation. By choosing one state or the other, the system breaks its spatial symmetry, just like by decreasing σ_1 or σ_2 below their critical values in the agent-resource system we see a sudden transition from an equal distribution of agents on the two resources to a definite bias towards one or the other.

Time-dependent solution

So far we have described results derived from the time-independent simplification of the model [9]. In order to extend our simulations to account for time-dependent behaviour we began by studying the evolution of a system with two resources and a number of agents $10 < N < 50$. We used the deterministic equation (see Figure 1) with the linear pay-offs as in (1). Each pay-off is described by a linear equation in the percentage of occupancy f of resource 1 (bear in mind that the system is closed, thus $f_2 = 1 - f$) as,

$$G_1 = af + b \quad \text{and} \quad G_2 = cf + d \quad (3)$$

If we start from an initial distribution f , which will depend on σ_1 and σ_2 , the less-dominant resource (G_2) mutates (increases) its slope and intercept proportionally to $\Delta f = f - 0.5$. In order to constrain the system we postulate that these increases




Figure 4: Effect of the uncertainty parameters σ_1 and σ_2 on a bistable system.

are equally matched by the decreases in slope and intercept for resource 1,

$$G_2 = (c + \Delta c)f + (d + \Delta d) \quad \text{and} \quad G_1 = (a - \Delta c)f + (b - \Delta d) \quad (4)$$

where,

$$\Delta c = \gamma \Delta f + \delta \quad \text{and} \quad \Delta d = \alpha \Delta f + \beta \quad (5)$$

each has two contributions: one depending on how badly they are loosing to the competing resource and another random component introducing noise. For simplicity we have adopted $\alpha = \gamma$ and $\delta = \beta$.

The competing process goes as follows: the deterministic equation gives an initial distribution f that allows each resource to calculate how much they have to mutate their pay-offs to become more attractive to the agents. The new pay-offs are re-introduced, together with σ_1 and σ_2 , to calculate the new probability ρ which is used to solve the deterministic and the fluctuations equations simultaneously.

As in biology, the rate at which mutations happen is fundamental to achieving an evolutionary improvement. In this simple case we have observed that after a mutation, and whichever initial configuration we start from, equilibrium is always reached after 2 units of time. This is of the order of the relaxation time of the evolution equations. In Figure 5 we show the evolution of a system with two resources with pay-offs as in (1), with $\sigma_1 = \sigma_2 \approx 0$ and a relaxation of 2 units of time. The simulation starts from resource 1 having its maximum of probability centered on 73% of the number of agents. At $t = 5$ resource 2 has pulled back and is winning more than half of the agents, while at $t = 10$ the situation goes back to a balance (50%). After this point the simulation is dominated by the noise introduced by the β and δ values (see (5)).




Figure 5: Evolution in a competitive system with linear pay-offs

4 Conclusions

We have studied a model for market-like agent-resource systems, whose formulation is based on one-step Markov processes and the large system-size expansion of the master equation due to Van Kampen [12] [13]. Our formulation enables a systematic treatment of fluctuations to be carried out. A deterministic equation governing the dynamics of the system arises as the lowest order contribution in the expansion, and coincides with the equation obtained in the mean-field approach [11]. The next order term gives the main contribution of the fluctuations in the form of a linear Fokker-Planck equation. The probability distribution describing the dynamics of the system is therefore a Gaussian distribution to this order in the expansion. Higher order terms are included to provide non-linear corrections to the FPE. the lowest order non-linear corrections represent fluctuations due to individual agents in the system. Higher order corrections are crucial when the number of agents is relatively small and mean-field theory inadequate.

To test the approximation in the case of our agent-resource system, we have taken a system with two resources and considered time-independent states. Taking numerical values as in ref. [11] shows full agreement between our exact theoretical results and the corresponding Monte Carlo simulations performed in ref. [11]. Sensitivity to accuracy of the information available to agents was also studied and the main observation is that higher uncertainty leads to the suppressing of non-linear noise effects. In view of the results obtained we conclude that the approximation works generally well and can therefore be reliably used for time-dependent solutions.

We have modelled time evolution by making two resources with linear pay-offs compete for the agents; after an initial period of instability the system adopts a *tit-for-tat* cycle where one resource dominates the other only to give way to the competing one after a fixed relaxation time interval. The time-dependent solutions for a small number of agents (for which first and higher order approximations are required) have been studied [14]. In that study, more realistic descriptions of the pay-offs in terms of systems' measurable properties and more sophisticated evolution mechanisms are described.

We should note, however, that Van Kampen's approximation scheme is not

suited for the treatment of fluctuations in situations involving instabilities or critical behaviour. In other words, the system size expansion is valid only when there is one globally stable macroscopic solution (such as a simple competitive system) or in the immediate vicinity of a locally stable solution. In general, however, a system with multiple minima requires a different treatment of fluctuations near instability points. This important issue will be addressed in a separate work.

References

- [1] M Gell and I Adjali, “Markets in Open Telecommunication Service Systems”, *Telematics and Informatics*, Vol 10, No 2, p. 131, 1993.
- [2] L Kleinrock, “Performance Models for Distributed Systems”, in *Teletraffic Analysis and Computer Performance Evaluation*, edited by O J Boxma, J W Cohen and H C Tijms, Elsevier Science Publishers BV, North Holland, p. 1, 1986.
- [3] J Power, “Distributed Systems and Self-organisation”, *Proceedings of the ACM 18th Annual Computer Conference*, pp. 379-384, 1990.
- [4] H Haken, *Information and Self-organisation - A Macroscopic Approach to Complex Systems*, Sringer-Verlag, 1988.
- [5] E Callen and D Shapero, “A Theory of Social Imitation”, *Physics Today*, Vol 27, p. 23, 1974.
- [6] W B Arthur, “Self-reinforcing Mechanisms in Economics”, in *The Economy as an Evolving Complex System*, Santa Fe Institute Studies in the Sciences of Complexity, Addison-Wesley, 1988.
- [7] P Bak and K Chen, “Self-organized Criticality”, *Scientific American*, p. 26, January 1991.
- [8] S Olafsson and M Gell, “Application of an Evolutionary Model to Telecommunication Services”, *European Transactions on Telecommunications*, Vol 4, p. 69, 1993.
- [9] I Adjali, M Gell and T Lunn, “Fluctuations in a Decentralised Agent-Resource System”, *Physical Review E* 49, p. 3833, 1994.
- [10] B A Huberman and T Hogg, in *The Ecology of Computation*, edited by Huberman, North Holland, Amsterdam, 1988.
- [11] J O Kephart, T Hogg and B A Huberman, “Dynamics of computational ecosystems”, *Physical Review A* 40, p. 404, 1989.
- [12] N G Van Kampen, “A power series expansion of the master equation”, *Can J Phys*, Vol 39, p. 551, 1961.
- [13] N G Van Kampen, in *Stochastic Processes in Physics and Chemistry*, North Holland Publishing Company, 1981.
- [14] J L Fernández-Villacañás, M Gell and I Adjali, “Evolution in a Self-Organising Agent-Resource System”, in *Proceedings of the IEEE Conference on Evolutionary Computation (IEEE WCCI 94)*, Orlando, Florida, June 27-29 1994.